




High-Fidelity Aerospace Simulations in the Exascale Era

Dimitri Mavriplis
University of Wyoming


Overview

- Overview of past progress in CFD for aerospace applications
- Assessment of future progress in Simulation Capabilities for aerospace applications
- Importance of contributing technologies
- Can we get there from here on our current path ?

CFD2030 Roadmap




Progress in Simulation Capabilities




- How do we measure Simulation Capability
- Suboptimal may be expected
- Is Super-Linear even possible ?

Progress in Simulation Capabilities




Improvements in Algorithms Relative to Moore's Law



The relative gains in some algorithms for the solution of an electrostatic potential equation on a uniform cubic grid compared to improvements in the hardware (Moore's Law).



- How do we measure Simulation Capability
- Suboptimal may be expected
- Is Super-Linear even possible ?
 - Combination of algorithmic and hardware advances

Progress in Simulation Capabilities





- Moore's Law is nominal : to be expected

Progress in Simulation Capabilities



- Moore's Law is nominal : to be expected
- RANS Plateau probably looks like this

Progress in Simulation Capabilities



- Can we re-invigorate progress through increased investment in fundamental disciplines ?
 - Will be required to meet CFD2030 Roadmap

Circa 1987

AIAA'87



AIAA-87-0452
Improvements to the Aircraft Euler Method
A. Jameson and T.J. Baker
Princeton University, Princeton, N.J.

AIAA 25th Aerospace Sciences Meeting
January 12-15, 1987/Reno, Nevada

For permission to copy or republish, contact the American Institute of Aeronautics and Astronautics
1633 Broadway, New York, NY 10019



- 2nd “Airplane” Paper
- Delaunay triangulation
- Unstructured mesh Euler solver
 - JST Scheme
 - Explicit Runge-Kutta
 - Implicit residual smoothing
 - Enthalpy damping

1987 Jameson Airplane Paper



- Unstructured tetrahedral mesh
 - **35,370 points, 181,959 tetrahedra**
 - Mesh generation: 15 minutes
 - No mention of geometry issues
 - Flow solver : 1 hour on 1 processor of CRAY-XMP
 - Vectorized, later parallelized for CRAY-XMP/YMP


1987 Jameson Airplane Paper



- Unstructured tetrahedral mesh
 - **35,370 points, 181,959 tetrahedra**
 - Mesh generation: 15 minutes
 - No mention of geometry issues
 - Flow solver : 1 hour on 1 processor of CRAY-XMP
 - Vectorized, later parallelized for CRAY-XMP/YMP

Circa 1999 (12 Years later)

Application Domain:
Computational Aerodynamics



SC'99

Gordon Bell Prize Finalist Talk

- **PETSc-FUN3D wins**
1999 Gordon Bell prize
Anderson, Keyes and Gropp



AIAA 99-0537

LARGE-SCALE PARALLEL
UNSTRUCTURED MESH
COMPUTATIONS FOR 3D HIGH-LIFT
ANALYSIS

D. J. Mavriplis

Institute for Computer Applications in Science and
Engineering


MS 403, NASA Langley Research Center
Hampton, VA 23681-0001

S. Pirzadeh

Configuration Aerodynamics Branch
MS 499, NASA Langley Research Center
Hampton, VA 23681-0001



37th AIAA Aerospace Sciences Meeting
January 11-14 1999, Reno NV

1999 High Lift Paper





- Coarse Mesh: 3 million points
- Fine mesh: 25 million points
- RANS simulation on up to 1500 CRAY-T3E processors
 - c/o Rob Vermeland

1999 High Lift Paper



- Coarse Mesh: 3 million points
- Fine mesh: 25 million points
- RANS simulation on up to 1500 CRAY-T3E processors
 - c/o Rob Vermeland

1999 High Lift Paper



- Reasonable agreement with experimental force data
- Easier take-off configuration

Evaluating Progress 1987-1999

- 12 years of Moore's Law: $2^8 = 256$
 - 256 X 35,370 points = 9M points
- Actual Increase in Capability
 - 25M points/35,370 = **700**
 - Euler to RANS
 - 10 to 50 times more computational requirements
 - Actual increase > **10,000**
 - Equivalent to 20 years of Moore's Law
 - Enabled by combination of hardware and algorithmic advances
 - Advancing layers mesh generation
 - Implicit/multigrid solvers
 - SA Turbulence model 1992
 - MPI as standard for parallel computing

1999 – 2017 (HLPW3)


Case 1 Grid Families



Series	Type	Points (M)	Cells (M)	Developer	Tool
A	Overset	24, 65, 189, 565	23, 64, 185, 554	NASA Ames	ANSA+CGT
B1	Tetrahedral	8, 26, 70, 206	48, 157, 416, 1228	Pointwise	Pointwise
B2	Mixed (prism dominant)	8, 26, 70, 206	22, 65, 170, 541	Pointwise	Pointwise
B3	Mixed (hex dominant)	8, 27, 71, 208	18, 48, 119, 397	Pointwise	Pointwise
C	Structured point-matched	10, 77, 338	8, 68, 311	GridPro	GridPro

Blue=Coarse
Green=Medium
Black=Fine
Red=X-Fine

c/o T. Sclafani







- 1999 Fine Grid is ~ 2017 Medium Grid
- Finest 2017 grid ~ 10X finest 1999 grid
 - Only 5 years equivalent of Moore's Law over actual 18 year period

HL-CRM coefficient of variation (Case 1a, “fine” grid)

Case	Unstr “F” grid sizes	Cv for lift	Cv for drag
HiLiftPW-1, alpha=13	31-162 M points	0.014	0.021
HiLiftPW-1, alpha=28		0.017	0.020
HiLiftPW-2, alpha=7	73-177 M points	0.025	0.020
HiLiftPW-2, alpha=16		0.023	0.028
HiLiftPW-3, alpha=8	70-189 M points	0.022	0.020
HiLiftPW-3, alpha=16		0.023	0.023

**No noteworthy decrease in Cv over the course of the 3 workshops
(but the “F” grids have not gotten much finer, either!)**



HL-CRM C_L statistics, Case 1a, alpha=16 deg.



c/o C. Rumsey


Slow Growth in Grid Size

- Grid resolution demonstrated to reduce scatter
- Why slow growth in grid size ?
 - Good enough for engineering
 - Lack of computing resources
 - Mesh generation does not scale
 - Flow solver does not scale
 - Flow solution not optimal $O(N)$
 - $O(N^2)$: 8x finer grid cost 64 times more to solve
- **Static technology** : Inability to leverage Moore's law (and new hardware)




Silver Linings from HLPW3

- Possible renaissance in new technology
 - GMGW Workshop
 - Anisotropic mesh adaptation
 - Transition prescription/prediction
 - New discretizations
 - Lattice Boltzmann, SUPG ($p=1$, $p=2$)
 - Strong solvers
 - More robust
 - but not scalable/optimal



Silver Linings from HLPW3


- Possible renaissance in new technology
 - GMGW Workshop
 - Anisotropic mesh adaptation
 - Transition prescription/prediction
 - New discretizations
 - Lattice Boltzmann, SUPG ($p=1$, $p=2$)
 - Strong solvers
 - More robust
 - but not scalable/optimal




c/o F. Alauzet

Silver Linings from HLPW3

- Possible renaissance in new technology
 - GMGW Workshop
 - Anisotropic mesh adaptation
 - Transition prescription/prediction
 - New discretizations
 - Lattice Boltzmann, SUPG ($p=1$, $p=2$)
 - Strong solvers
 - More robust
 - but not scalable/optimal



New Advocacy



- Capability stagnation has led to renewed advocacy
 - Importance of algorithmic developments in enabling capability advances
 - Numerical methods
 - Computer science
 - Physical modeling

Old Advocacy

18th AIAA Computational Fluid Dynamics Conference
25 - 28 June 2007, Miami, FL

AIAA 2007-4084

Petaflops Opportunities for the NASA Fundamental Aeronautics Program

Dimitri J. Mavriplis *
David Darmofal †
David Keyes ‡
Mark Turner §

The premise of this paper is the observation that the engineering community in general, and the NASA aeronautics program in particular, have not been active participants in the renewed interest in high performance computing at the national level. Advocacy for high performance computing has increasingly been taken up by the science community with the argument that computational methods are becoming a third pillar of scientific discovery alongside theory and experiment. Computational engineering, on the other hand, has continually been relegated to a set of mature software tools which run on commodity hardware, with the notion that engineering problems are not complex enough to warrant the deployment of state-of-the-art hardware on such a vast scale. We argue that engineering practices can benefit equally from an aggressive program in high performance computational methods, and that these problems are at least as important as science problems, particularly with regards to any national competitiveness agenda. Because NASA aeronautics has historically been a principal driver of computational engineering research and development, the current situation represents an opportunity for the NASA aeronautics program to resume its role as a leading advocate for high performance computational engineering at the national level. We outline a sample set of Grand Challenge problems which are used to illustrate the potential benefits a reinvigorated program could produce, and use these examples to identify critical barriers to progress and required areas of investment. We conclude by noting that other communities have spent significant efforts in formulating the case for increased investment in high performance computing activities, and that a similar roadmap will be required for the engineering community.

- 10 years of Moore's Law: 100x
- Evaluate capability growth since 2007
- Compare to other IT "technology" growth since 2007

Old Advocacy

18th AIAA Computational Fluid Dynamics Conference
25 - 28 June 2007, Miami, FL

AIAA 2007-4084

Petaflops Opportunities for the NASA Fundamental Aeronautics Program

Dimitri Mavriplis (University of Wyoming)

David Darmofal (MIT)

David Keyes (Columbia University)

Mark Turner (University of Cincinnati)

18th AIAA Computational Fluids Dynamics Conference
25-28 June, 2007, Miami Florida

- 10 years of Moore's Law: 100x
- Evaluate capability growth since 2007
- Compare to other IT "technology" growth since 2007

Old Advocacy

18th AIAA Computational Fluid Dynamics Conference
25 - 28 June 2007, Miami, FL

AIAA 2007-4084

Petaflops Opportunities for the NASA Fundamental Aeronautics Program

Dimitri Mavriplis (University of Wyoming)

David Darmofal (MIT)

David Keyes (Columbia University)

Mark Turner (University of Cincinnati)

18th AIAA Computational Fluids Dyna
25-28 June, 2007, Miami

- 10 years of Moore's Law: 100x
- Evaluate capability growth since 2007
- Compare to other IT "technology" growth since 2007



1st iPhone introduced 2007

Old Advocacy

18th AIAA Computational Fluid Dynamics Conference
25 - 28 June 2007, Miami, FL

AIAA 2007-4084

Petaflops Opportunities for the NASA Fundamental Aeronautics Program

Dimitri Mavriplis (University of Wyoming)
David Darmofal (MIT)
David Keyes (Columbia University)
Mark Turner (University of Cincinnati)

18th AIAA Computational Fluids Dyna
25-28 June, 2007, Miami F

- 10 years of Moore's Law: 100x
- Evaluate capability growth since 2007
- Compare to other IT "technology" growth since 2007





1st iPhone introduced 2007



Emergence of AI/ML

Science Runs on Red Storm




- SEAM (Spectral Element global Atmospheric Model) Simulation of the breakdown of the polar vortex used to study the enhanced transport of polar trapped air to mid latitudes.
- Record setting 20 day simulation, 7200 cpus for 36 hours. 1B grid points (3000x1500x300), 300K timesteps, 1TB of output.
- Spectral elements replace spherical harmonics in horizontal directions
- High order ($p=8$) finite element method with efficient Gauss-Lobatto quadrature used to invert the mass matrix.
- Two dimensional domain decomposition leads to excellent parallel performance.

c/o Mark Taylor, Sandia National Laboratories


18th AIAA Computational Fluids Dynamics Conference
25-28 June, 2007, Miami Florida

SEAM on Red Storm and BG/L


Parallel Scalability



Max: 5TF



Parallel Scalability




Max: 4TF




Performance of 4 fixed problem sizes, on up to 6K CPUs. The annotation gives the mean grid spacing at the equator (in km) and the number of vertical levels used for each problem.

Current State-of-the-Art HPC



Gordon Bell 2015: Earth Mantle simulation 1.6M cores



Wyoming in-house DG code (512³ mesh @p=4) using up to 1 million MPI ranks (2 per core) on Mira at Argonne




Gordon Bell 2016:10M-Core Scalable Fully-Implicit Solver for Nonhydrostatic Atmospheric Dynamics

Factor 100 over 10 years held up at high end of HPC

NASA Computational Environment

- Columbia processes mostly $O(100)$ cpu jobs
- 2048 sub-system occupied with 512 jobs
- Few benchmarks above 512 cpus
- Some 2048 benchmarks (production ?)



Aerospace Computational Environment

- Does NASA Pleiades process mostly 10,000 core jobs ?
 - x100 from 2007 Columbia
- DoD HPCMP Machines offer large job allocations
- In general production size jobs have not grown X100
 - Stagnation of grid sizes at HLPW3
- Computing is more ubiquitous
 - Larger number of jobs possible
- Heterogeneous architectures still not mainstream



Selected Grand Challenges

- Digital Flight
 - Static (and dynamic) aerodynamic data-base generation using high-fidelity simulations
 - Time-dependent servo-aero-elastic maneuvering aircraft simulations
- Transient Full Turbofan Simulation
- New frontiers in multidisciplinary optimization
 - Time dependent MDO
 - MDO under uncertainty
- Examples only (not all inclusive)
 - e.g. Aeroacoustics not mentioned





Design Optimization Challenges

- Unsteady Multidisciplinary Design Optimization:
 - Adjoint methods require backwards integration in time
 - Requires entire time-dependent solution set to be stored (to disk)
- Design under uncertainty
 - Ensemble averages for uncertainty estimation
 - Stochastic methods

Computational Requirements

- One analysis cycle
 - 100 million grid points, one revolution
 - ~30 hours on 100 cpus
- One design cycle (twice cost of analysis)
 - Forward time dependent simulation
 - Backwards time dependent adjoint solution
- 50 to 100 design cycles
- 30 to 60 hours on 10,000 cpus

Time Dependent Multi-Disciplinary Optimization




Rigid and Flexible blades

Unsteady adjoint optimization in FUN3D (2010-2013)

Far-field acoustic optimization of flexible rotor (Fabiano and Mavriplis (2016))

10m points on 4096 cores for 20 hours (largest to date)





Computational Requirements


- From M. D. Salas (2006): *Digital Flight: The last CFD Aeronautical Grand Challenge*
 - 60 seconds of flight = 1.5 days on 512 cpus
 - NASA codes, 50 million grid points, 50Hz time stepping
- Easily add:
 - Order of magnitude in grid resolution
 - Order of magnitude in time resolution
 - Multidisciplinary:
 - Structures, Heating, Flight control system
 - Overnight turnaround on 10,000 cpus

Substantial Advances in Digital Flight

CREATE-AV





- Leveraged dynamic overset, AMR, higher order, multidisciplinary
- Digital flight for rotorcraft even more challenging



Wyoming Wind Energy Simulations

- Highly interdisciplinary
 - Aero, structures, controls, atmospheric turbulence
- Technology enablers
 - Unstructured mesh solvers
 - Dynamic adaptive meshing
 - Dynamic overset meshes
 - High-order (DG) off-body
 - LES modeling
 - Atmospheric boundary layer modeling
- Exascale problem
 - 10 orders of magnitude range of scales




48 Wind Turbine Simulation

1.5B dof's

22K cores

1.28 revs in
12 hours

Scales:
10 km to
7microns



UNIVERSITY OF WYOMING

UNIVERSITY OF WYOMING

	Turbine Count	Efficiency	Revs	Near-Body Cores	Off-Body Cores	Total Cores
25 mesh bodies →	6	1.0000	1.374	2,088	720	2,808
	12	0.9874	1.360	4,176	1,440	5,616
	24	0.9682	1.331	8,352	2,880	11,232
	48	0.9333	1.283	16,704	5,760	22,464
385 mesh bodies →	96	0.8686	1.194	33,408	11,520	44,928

Good weak scaling from 6 – 96 turbines

Summary


- Difficult Problems

- Our ability (at UW) to simulate C_{Lmax} is about as good as it was 15 years ago: **Not very good**
- C_{Lmax} for HL will require advances in :
 - Geometry modeling
 - Grid generation/adaptivity
 - Solver technology
 - Physical modeling
 - Transition, Turbulence modeling (RANS, LES)





- Substantial advances can be made simultaneously in other areas

- Multidisciplinary simulations
- Optimization technology
- Uncertainty quantification
- Automated/Robust Data-base fill-in




Conclusions

- Advances in complementary fundamental disciplines are required to **simply keep pace** with Moore's Law
 - Offer the possibility of outperforming Moore's Law
- Even more true with increasing computational power
 - Asymptotically arguments most powerful at large scale
- Required to meet the CFD 2030 Vision



The relative gains in some algorithms for the solution of an electrostatic potential equation on a uniform cubic grid compared to improvements in the hardware (Moore's Law).



Acknowledgements


- Mike Brazell, Behzad Aharabi, Zhi Yang, Andrew Kirby, HLPW3 Committee
- University of Wyoming Advanced Research Computing Center (ARCC)
- NCAR-Wyoming Supercomputer Alliance



EPSCoR



Sea-Based Aviation



T³ Project



Univ of Maryland